

Constrained Signaling for Welfare and Revenue Maximization

Shaddin Dughmi * Nicole Immorlica † Aaron Roth ‡

February 19, 2013

Abstract

We consider the problem of an auctioneer who faces the task of selling a good (drawn from a known distribution) to a set of buyers, when the seller does not have the capacity to describe to the buyers the exact identity of the good that he is selling. Instead, he must come up with a constrained *signalling scheme*: a (non injective) mapping from goods to signals, that satisfies the constraints of his setting. For example, the auctioneer may be able to communicate only a bounded length message for each good (equivalently, he may be constrained to using only a fixed number of signals in total). The auctioneer may also face additional exogenously imposed constraints on the signaling scheme – for example, he might be legally constrained to truthfully advertise the item being sold. Each candidate signaling scheme induces an incomplete-information game among the buyers, and the goal of the mechanism designer is to choose the signaling scheme that optimizes either welfare or revenue. In this paper, we give algorithms for computing constrained signaling schemes, as well as hardness results, for both of these objectives for a variety of constrained signaling problems.

1 Introduction

At a cafe in *Portlandia* [ABK11], customers about to order the chicken ask the waitress for more details regarding its source. She informs them that the chicken is a heritage breed, woodland raised, and has been fed a diet of sheep’s milk, soy, and hazelnuts, and she assures them that the chicken is indeed local, free range, and so on. Ah but, the customers ask, is that USDA Organic, or Oregon Organic, or Portland Organic?

Organic certifications are a means by which a seller can communicate parameters of a product to potential buyers. The certification system creates a simplified and practical, yet sufficiently expressive, set of signals as the basis for this communication.¹ Signaling is common throughout complex markets. The USDA classifies meat into a small number of discrete grades. Graduating high-school students signal their potential to employers and colleges through SAT scores and transcripts. Targeted advertising sales in both online and offline media describe viewers to ad buyers through a fixed and (comparatively) small set of demographic information.²

*University of Southern California

†Microsoft Research

‡University of Pennsylvania. This work was partially supported by an NSF CAREER Grant and NSF Grant CCF-1101389

¹Unfortunately for the waitress in *Portlandia*, the language was not sufficiently expressive for her customers. They decided to drive to the farm to visit the chicken’s home.

²See Milgrom [Mil10] and Levin and Milgrom [LM10] for further discussion of signaling in various markets including online advertising, wheat sales, diamond sales, and spectrum auctions.

A *signaling scheme* for a market maps each potential item to a signal. Buyers, upon viewing the signal, use their knowledge of the signaling scheme to refine their beliefs about the identity of the item for sale and update their value accordingly. For example, a seller of skinny jeans participating in an online advertising auction might have a high value for viewers that are hipsters and no value for retirees or rappers. Thus a signaling scheme that indicates age and geographical location will enable the advertiser to distinguish 20-somethings in Brooklyn from 70-somethings in Fort Lauderdale and thus separate retirees from hipsters, but conflate hipsters and rappers. If viewers are equally likely to be hipsters, retirees, or rappers, and the advertiser learns the viewer is a 20-something from Brooklyn, then he updates his expected value for the viewer to be half that of his value for a hipster. The choice of signaling scheme induces a game amongst buyers: the buyers find themselves competing in an auction without knowing what the good for sale is, but rather only knowing that it is some good that induced a signal s from a known signaling scheme. Thus, the choice of signaling scheme by the auctioneer corresponds to a particular refinement of the information-set presented to the buyers, and different choices can result in auctions that result in different degrees of welfare or revenue.

Practical settings impose a variety of objectives and constraints on this signaling game. First and foremost, for a signaling scheme to be useful as a means of communication, it must satisfy certain informational constraints. In particular, in many markets, the *amount* of information that can be conveyed is highly constrained. In online advertising auctions, the sheer volume of sales and diversity of viewers make it impractical to communicate precise details of every viewer. In addition, various reputational and legal constraints may restrict a seller from announcing certain signals for certain items: e.g., organic products must pass a certification process from the corresponding agency to be sold under that label. The choice of sale mechanism is also up for debate. Sellers typically wish to maximize particular objectives, like welfare or revenue, and so would like to choose a signaling scheme and corresponding mechanism that achieves that goal. If this were the extent of the problem, a seller could choose a revenue or welfare optimal mechanism, subject to the constraints of the setting. However, in multi-item settings with complex valuations, these mechanisms are often complicated. Furthermore, due to historical reasons, legal reasons, or buyer anonymity, sellers often face additional constraints on feasible mechanisms. In online advertising auctions, for example, the industry standard sells a single ad slot with a second price auction. Alternate mechanisms may enjoy better theoretical properties, but would receive significant backlash from the buyer base.

In this paper, we study welfare and revenue optimal signaling schemes in which the seller is constrained in both the feasible schemes and the sale mechanism. We consider the natural constraint of requiring our signaling schemes to have bounded *cardinality*. This is equivalent to having a bound on the communication that is allowed when transmitting a signal (it takes at most $\log k$ bits to transmit a signal that is one among a set of k possible signals). We sometimes also consider other, exogenously imposed constraints. In bipartite signaling, we are given a subset of feasible signals, and each item is allowed to be mapped to any one of a subset of these signals. For example, a signaling scheme for online advertising auctions may consist of all potential tuples of attributes: age, location, ethnicity, gender, and income, say. A viewer can be matched to any tuple that matches his true characteristics. So a hipster who is a white 20-something Brooklyn boy with no income can be mapped to the “20-somethings in Brooklyn” signal, in which case he might be conflated with a rapper, or the “person with no income” signal, in which case he might be conflated with a retiree. Cardinality-constrained signaling is a special-case of bipartite signaling in which any item can be mapped to any signal, but the number of signals is constrained. This special case

makes sense in settings where items can be grouped arbitrarily. We prove positive and negative results for bipartite and cardinality constrained signaling, in a setting with a discrete set of items and signals.

Finally, consider communication-bounded signaling, in which the constraint is the amount of communication which can be used to send the signal³. Our hipster may visit a bar and be faced with the decision of whether or not to order an “Old Chub” Ale. This micro-brew is a complex chemical mixture, with a consistency and taste that is impossible to fully describe in a reasonable amount of time. Instead, the brewery chooses a one sentence description⁴ which fails to capture every detail, but is the best description of the item given the space constraints. We give general signaling schemes for items which are described as vectors in d dimensional real valued space, for bidders who have geometric valuation functions for the goods, which can be communicated using $\ll d$ bits. In this setting, we are able to give strong results in the unknown valuation setting, in which valuations can be drawn from an arbitrarily correlated prior, sometimes even when it is *unknown* to the mechanism choosing the signalling scheme.

We impose several additional assumptions on our setting. First, as mentioned above, we assume the seller is constrained to use a second-price auction.⁵ Second, the choice of signaling scheme is clearly influenced by what the seller knows about buyers’ values. If all buyers that value hipsters also value rappers, then there is no loss in mapping these two viewer types to the same signal. As is standard, we model values as being drawn from a common prior and allow the seller to choose a signaling scheme based on this prior. We call this the *unknown valuation* setting, and study the properties of the resulting second-price auction in the induced incomplete information game. We are able to give strong results in this setting for geometric valuation functions, in the bounded communication regime – sometimes even when the prior itself is unknown to the mechanism designer. Unfortunately, some of our problems prove fairly difficult in this setting, and thus we also study the somewhat less natural *known valuation* setting in which the seller knows the set of values of the buyers precisely. This induces a complete information game. Our motivation for studying this setting is two-fold: for one, we hope that it will lead to mechanisms for the unknown valuations case, and indeed some of our results carry through when the amount of uncertainty in the unknown valuations case is bounded in some sense.

We also note that our mechanisms for known valuations treat buyers anonymously. Thus these mechanisms might be applicable to large markets in which the seller can be fairly certain of the set of values but not of the identity of the different buyers. Specifically, consider settings where a seller expects customers to belong to one of a number of demographics, or types, with one customer from each type arriving with high probability. Then with high probability, several customers of each type arrive. In this case, it is easy to see that the revenue and welfare of a signaling scheme are equal, with high probability, to the welfare of the signaling scheme for customers with known valuations equal to the types. Therefore, if an approximately welfare-maximizing scheme can be computed in polynomial time in the known-valuations variant of the problem, a revenue and welfare maximizing signaling scheme can be computed in the unknown valuation setting as well.

³This is of course the same as cardinality-constrained signaling, but the change in perspective of viewing the constraint as being the description length of each signal can be useful.

⁴“This jaw-dropping Scottish strong ale (8% ABV) is brewed with bodacious amounts of malted barley and specialty grains, and a dash of beechwood-smoked malt.”

⁵Alternatively, we could state our results in terms of posted-price mechanisms; this does not significantly alter the underlying optimization problem and so most of our results apply to this alternate assumption as well.

Our Contributions We first observe that welfare maximization in constrained signaling with unknown valuations is a convex function maximization problem. Thus there is always a deterministic signaling scheme which maximizes expected welfare. This significantly simplifies our analysis, as we only have to search over deterministic signaling schemes. Furthermore, we observe that revenue and welfare are intimately related: the revenue is at most the welfare after excluding an arbitrary player. This allows us to extend some of our welfare results to the revenue setting.

We next study the bipartite signaling setting. In the known valuations case, we are able to reduce the welfare problem to submodular function maximization subject to a matroid constraint, thus implying a $(1 - 1/e)$ -approximation for welfare maximization. Starting from our approximately welfare-maximizing signaling scheme, and using the connection between revenue and welfare described above as well as careful “mixing” of different signals, we devise a constant approximation for revenue maximization in the cardinality-constrained signaling setting. We further show how to extend some of our results to the unknown valuation case when the prior has constant-size support. We couple our results with a hardness result: even in the cardinality-constrained signaling setting with known valuations, it is NP-hard to approximate either objective to a factor better than $(1 - 1/e)$ via a reduction from max-cover.

Finally, we study communication-constrained signaling. We consider settings in which goods are points in a very high dimensional space $\Omega = \mathbb{R}^d$, where d is much too large to communicate the entire vectors. In these settings, we consider agents who have geometric utility functions, defined in terms of either the distance or the angle from the realized good to some target good or set of goods. We use techniques from no regret learning and metric embeddings to give highly space efficient signaling schemes that achieve nearly the optimal welfare (compared even to the optimal unconstrained signaling scheme), for the unknown valuation case, and in one instance, even against *adversarially chosen* goods, and without knowledge of the prior distribution over valuations.

Related Work The study of markets with information asymmetries between sellers and buyers was first introduced by Akerlof [Ake70]. Since then, a rich literature has examined the effects of information revelation, i.e. signaling, in markets and auctions. Most notably, the “Linkage principle” of Milgrom and Weber [MW82] shows that, under some conditions, a seller always increases his revenue by signalling more information about the good for sale. However the linkage principle requires fairly strong assumptions regarding the joint distribution of player valuations for the good. In particular, it rarely holds in settings where players come from different demographics with negatively correlated values for the good. A review of the literature studying the limits of the linkage principle can be found in Emek et al [EFG⁺12]. The non-applicability of the linkage principle to the settings we consider goes even beyond these limitations — even when full transparency is optimal for an auctioneer, the presence of constraints on the amount and nature of information revealed by our signalling schemes introduces intricate tradeoffs in choosing *which* information to reveal. Quantifying those tradeoffs inevitably requires examination of these settings with an optimization lens, as we do here.

Our work is inspired by, and builds on, recent papers that examine optimal signalling schemes in *unconstrained settings*. Specifically, Emek et al [EFG⁺12] and Miltersen and Sheffet [MS12] examine signalling for revenue maximization in a second price auction, where no constraints are placed on the number or nature of signals. In such settings, full information revelation is optimal for an auctioneer interested in maximizing welfare, and both works show that a revenue-maximizing scheme can be computed efficiently when player valuations are known. Emek et al [EFG⁺12] also

obtain partial results for revenue maximization when player valuations are drawn from a Bayesian prior. Our results can be thought of as the extension of these works to settings where social, legal, or practical constraints are placed on the auctioneer’s signaling policy.

We also mention some results that are related to ours in the techniques used. First, the problem of computing a revenue-maximizing bundling of goods considered in Ghosh et al. [GNS07] is both conceptually and technically similar to the optimization problem we face in our revenue-maximizing scheme for cardinality-constrained signaling (Section 4.2). The idea of “merging” signals in our setting is inspired by their algorithm, though the constraint on the number of signals in our setting poses additional technical hurdles. The results of Section 4 also heavily use techniques from the combinatorial auctions literature ([Fei06, DS06, DNS05]), as well as the submodular function maximization result of Vondrak [Von08].

In the geometric setting considered in Section 5, in which players value goods $y \in \mathbb{R}^d$ according to *inner product valuations*, our communication bounded signaling scheme applies the multiplicative weights framework of Arora, Hazan, and Kale [AHK12]. We use multiplicative weights in an unusual way, both as a no-regret learning algorithm and as a *compression scheme*. That is, we use both the fact that multiplicative weights can quickly learn a hypothesis \hat{y} that closely approximates a vector y with respect to a fixed number of inner product valuations, and the fact that the hypothesis \hat{y} can be concisely communicated by transmitting only the update operations of multiplicative weights, rather than the vector \hat{y} itself. We are not aware of multiplicative weights being used explicitly as a compression scheme elsewhere, although this is related to the use that no-regret algorithms have recently found in differential privacy [RR10, HR10, GRU12].

2 Preliminaries

2.1 Constrained Signalling Problems

We consider a setting in which there is a set Ω of possible items for sale, and a single item $j \in \Omega$ is drawn from a distribution $p \in \Delta_\Omega$. There is a set of n players, each of whom is equipped with a *valuation* $v_i : \Omega \rightarrow \mathbb{R}$ mapping items to the real numbers. The valuation profile (v_1, \dots, v_n) is drawn from a distribution \mathcal{D} . We assume that Ω , p , and \mathcal{D} are common knowledge, while each player’s valuation v_i is private to player i .

We assume that the realization of item j is ex-ante unknown to the players, but known to the auctioneer. We consider an auctioneer who first observes the drawn item j , and then announces a string s , known as a *signal*. The (possibly randomized) policy by which the auctioneer chooses his signal, which we refer to as a *signaling scheme*, is common knowledge. After players observe the signal s , which can be thought of as a random variable correlated with the realization of the item j , a second price auction for item j is run.

We adopt the perspective of an auctioneer seeking to optimize his choice of signaling scheme, with the goal of maximizing the expected welfare or revenue obtained in the second price auction. We focus on settings where, for practical, social, or legal reasons, the auctioneer is constrained in their choice of signaling scheme. We therefore define the class of constrained signaling problems, as follows.

Definition 2.1. *A constrained signaling problem is a family of instances, each of which is given by:*

- *A set Ω of items, and a distribution $p \in \Delta_\Omega$ over these items.*

- A set $[n]$ of players, where each player i is equipped by a private valuation $v_i : \Omega \rightarrow \mathbb{R}_+$. The tuple of valuations (v_1, \dots, v_n) is drawn from a common prior \mathcal{D} .
- A set of signals \mathcal{S} , and a set $\mathcal{F} \subseteq \mathcal{S}^\Omega$ of valid signaling maps.

When \mathcal{D} is a trivial prior, i.e. (v_1, \dots, v_n) are deterministic, we say our signaling problem has *known valuations*. A solution to a constrained signaling problem is a *valid signaling scheme*, defined as a distribution $x \in \Delta_{\mathcal{F}}$ over valid signaling maps \mathcal{F} . When x is a point distribution, we say the signaling scheme is *deterministic*. Given x and an item $j \in \Omega$, we use $x(j)$ to denote the random variable $f(j)$ for $f \sim x$. Moreover, given item $j \in \Omega$ and signal $s \in \mathcal{S}$, we abuse notation and use $x(j, s)$ to denote the probability that $f(j) = s$ for $f \sim x$. Similarly, given a signal $s \in \mathcal{S}$ we use $x(s)$ to denote the probability that $f(j) = s$ for $f \sim x$ and $j \sim p$; it is easy to see that $x(s) = \sum_{j \in \Omega} p_j x(j, s)$.

A signaling scheme x induces, for each signal s , a second-price auction where players have independent private values.⁶ Specifically, it is a dominant strategy for each player i to bid his value for item $j \sim p$ conditioned on signal s — namely

$$v_i|s, x := \mathbf{E}_{f \sim x, j \sim \mathcal{D}} [v_i(j) | f(j) = s] = \frac{\sum_{j \in \Omega} x(j, s) p_j v_i(j)}{x(s)}$$

The winning player for signal s is the player maximizing $v_i|s, x$. Parametrized by the valuation profile v , the resulting welfare of the auction, in expectation over all draws of the item j , is given by

$$welfare(x, v) := \sum_{s \in \mathcal{S}} x(s) \max_{i=1}^n v_i|s, x = \sum_{s \in \mathcal{S}} \max_{i=1}^n \sum_{j \in \Omega} x(j, s) p_j v_i(j).$$

Using $\hat{v}_i(j) := p_j v_i(j)$ to denote player i 's value for item j weighted by j 's probability, gives

$$welfare(x, v) = \sum_{s \in \mathcal{S}} \max_{i=1}^n \sum_{j \in \Omega} x(j, s) \hat{v}_i(j). \quad (1)$$

And similarly, the revenue in the second price auction is given by

$$rev(x, v) := \sum_{s \in \mathcal{S}} x(s) \max_{i=1}^n 2 v_i|s, x = \sum_{s \in \mathcal{S}} \max_{i=1}^n 2 \sum_{j \in \Omega} x(j, s) \hat{v}_i(j), \quad (2)$$

Where $\max 2$ denotes selecting the second largest value. In the case of unknown valuations, the expected welfare and revenue of the auction over draws of the players' valuations, which we denote by $welfare(x)$ and $rev(x)$, are the expectations of $welfare(x, v)$ and $rev(x, v)$, respectively, over $v \sim \mathcal{D}$.

2.2 Example Problems

Below we define some natural examples of constrained signaling problems.

⁶This crucially relies on the fact that our signaling schemes are “symmetric,” in the sense that the auctioneer announces the same information — i.e. signal — to all players. An “asymmetric” signaling scheme, which releases different information to different players, would induce a second price auction with interdependent values, which we avoid in this paper.

2.2.1 Bipartite and Cardinality-constrained Signaling

In a *bipartite signaling problem*, the set of valid signaling maps is represented explicitly as a bipartite graph with Ω on the left hand side, \mathcal{S} on the right hand side, and a set of edges $E \subseteq \Omega \times \mathcal{S}$, as well as an integer k . A signaling map $f \in \mathcal{S}^\Omega$ is valid if $(j, f(j)) \in E$ for each $j \in \Omega$, and moreover $|f(\Omega)| \leq k$. We refer to the edge set E as the *bipartite graph constraint*, which limits the compatible signals with each item, and the integer k as the *cardinality constraint*, which limits the total number of signals used. To ensure the existence of at least one feasible signaling map and simplify our results, we assume that in instances with $k < |\mathcal{S}|$, there exists a “no information” signal $s_0 \in \mathcal{S}$ such that $(j, s_0) \in E$ for all items $j \in \Omega$.

We assume the distribution p over items is given explicitly as a vector in $[0, 1]^\Omega$. In the known valuation variant of the problem, we also assume the valuation profile is given as a set of vectors $v_1, \dots, v_n \in \mathbb{R}^\Omega$. For the variant with unknown valuations, we assume for simplicity of presentation that the joint distribution \mathcal{D} over valuation profiles is given as a set of matrices $\{v^1, \dots, v^t\}$ where $v^k \in \mathbb{R}^{[n] \times \Omega}$, and associated probabilities q_1, \dots, q_t .

We refer to the special case of this problem where $E = \Omega \times \mathcal{S}$, i.e. where the only restriction on the signaling scheme is the number of allowed signals, as the *cardinality-constrained signaling problem*. Instances of cardinality-constrained signaling can equivalently be described by a set Ω and a number k of allowable signals (\mathcal{S} is implicitly assumed to be $\{1, \dots, k\}$ without loss of generality).

2.2.2 Communication Bounded Signaling

In many cases, a product y may be best described as a collection of d features where d is extremely large. In other words, goods may be points $y \in \mathbb{R}^d$ – natural valuation functions for such points take advantage of this Euclidean geometry, and might be represented as the dot product (i.e. the angle) between a good and some target point, or the distance between the good and some target point or subspace. However, because the dimension of the goods d is so large, it is not possible to exactly describe them to buyers, and it is instead necessary to employ a *communication bounded signaling scheme*. Such a scheme places a limit on the amount of communication that can be invested in transmitting the signal to the agents. Consider for example the sale of wine. Wine is a complicated, high dimensional product, but rather than being described by a full molecular description, it is sold with a description that is at most a paragraph in length. Informally, we say that a signalling scheme has b -bounded length of the total set of signals it can generate (over all possible products) is of size at most 2^b – note that any signal from such a set can be indexed by using at most b bits.

We say that a bidder i has *inner product valuations* if bidder i is parameterized by a vector $v_i \in \mathbb{R}^d$, and her value for a good $y \in \mathbb{R}^d$ is $\langle v_i, y \rangle$. In the known valuation setting, we assume that the vectors v_1, \dots, v_n are given explicitly. In the unknown valuation setting, we assume that there is some prior distribution \mathcal{P} from which the vectors v_i are drawn. Our algorithms do not require any explicit representation of \mathcal{P} , but only the ability to sample from it.

We say that a bidder has *subspace valuations* if bidder i is parameterized by k vectors $v_i^1, \dots, v_i^k \in \mathbb{R}^d$, and her value for a good $y \in \mathbb{R}^d$ is a function of $d(y, S)$, the Euclidean distance from the good y to the subspace spanned by v_i^1, \dots, v_i^k . Our algorithms for signalling with respect to subspace valuations work in the unknown valuation setting and yet are independent of the prior distribution on values, and so work for arbitrarily complicated priors, even without the ability to sample from them.

3 Structural Properties

3.1 Convexity and Determinism

Given a constrained signaling problem, we observe that welfare is a convex function of the marginal probabilities $x(j, s)$, and conclude that the welfare-maximizing constrained signaling scheme is deterministic. Formally, for every valid signaling map $f \in \mathcal{F} \subseteq \mathcal{S}^\Omega$, we associate a vector $x^f \in \{0, 1\}^{\Omega \times \mathcal{S}}$ where $x^f(j, s) = 1$ if and only if $f(j) = s$. A signaling scheme is then associated with a vector x in the convex hull of $\{x^f : f \in \mathcal{F}\}$, and has welfare and revenue as given in Equations 1 and 2, respectively. We observe $welfare(x)$ is a convex function of x , and invoke the following fact:

Fact 3.1. *Let \mathcal{P} be a polytope in Euclidean space, and let g be a convex function. The maximum of g over \mathcal{P} is attained at a vertex of \mathcal{P} .*

We conclude that there is a welfare-maximizing deterministic signaling scheme.

Lemma 3.2. *For any constrained signaling problem with unknown valuations, there is valid deterministic signaling scheme which maximizes expected welfare.*

3.2 Relating Revenue and Welfare

We now mention a useful, though elementary, upper bound on the optimal revenue achievable via a signaling scheme. This bound follows immediately from the fact that the revenue of a second price auction is at most the welfare of the same auction after excluding an arbitrary player, though we present a proof here for completeness.

Lemma 3.3. *Fix an arbitrary constrained signaling problem with unknown valuations. Let i' be an arbitrary player. The revenue of the revenue-optimal signaling scheme is at most the welfare of the welfare-optimal signaling scheme for all players other than i' .*

Proof. It suffices to show that $rev(x, v)$ (equation (2)) is at most $welfare(x, v_{-i'})$ (equation (1)).

$$\begin{aligned} rev(x, v) &= \sum_{s \in \mathcal{S}} \max_{i=1}^n \sum_{j \in \Omega} x(j, s) \widehat{v}_i(j) \\ &\leq \sum_{s \in \mathcal{S}} \max_{i \neq i'} \sum_{j \in \Omega} x(j, s) \widehat{v}_i(j) \\ &= welfare(x, v_{-i'}) \end{aligned}$$

□

4 Bipartite Signaling

In this section we present both positive and negative results for welfare and revenue maximization for bipartite signaling and the special case of cardinality-constrained signaling. In the known valuations case, we show the existence of a $1 - 1/e$ approximation algorithm for welfare maximization in bipartite signaling, and a constant approximation algorithm for revenue maximization in cardinality-constrained signaling. Our results extend to the unknown valuations case when the

Bayesian prior \mathcal{D} has constant-size support. Finally, we show that even in cardinality constrained signaling with known valuations, approximating welfare or revenue to a factor better than $1 - 1/e$ is NP -hard.

4.1 Welfare Maximization with Known Valuations

In the known valuations setting, we show the existence of a polynomial-time, $e/(e-1)$ -approximation algorithm for welfare maximization in bipartite signaling. Formally, we prove the following result.

Theorem 4.1. *For bipartite signaling with known valuations, there is a randomized, polynomial-time, $e/(e-1)$ -approximation algorithm for computing the welfare-maximizing signaling scheme.*

As a warmup, in Section 4.1.1, we show that the special case of bipartite signaling without the cardinality constraint reduces almost directly to combinatorial auctions with XOS valuations. We then extend the result to the general case in Sections 4.1.2 and 4.1.3 through a non-trivial reduction to submodular function maximization subject to a matroid constraint.

4.1.1 A Special Case: Without the Cardinality Constraint

We show that welfare maximization in bipartite signaling with known valuations and without a cardinality constraint — i.e. with $k = |\mathcal{S}|$ — reduces to welfare-maximization in combinatorial auctions⁷ with explicitly-represented XOS valuations,⁸ which can be approximated to within a factor of $\frac{e}{e-1}$ as shown by Donbuzinski and Schapira [DS06].

We are given a set of items Ω , a distribution $p \in \Delta_\Omega$ on items, a set of signals \mathcal{S} , a bipartite graph $E \subseteq \Omega \times \mathcal{S}$, and known valuations $v_1, \dots, v_n \in \mathbb{R}^\Omega$. As in Section 2, we use $\widehat{v}_i(j) = p_j v_i(j)$ as shorthand. A valid deterministic signaling scheme $f : \Omega \rightarrow \mathcal{S}$ partitions the items Ω among signals \mathcal{S} — we let $B_s = f^{-1}(s)$ denote the “bundle” of items mapped to signal $s \in \mathcal{S}$. Appealing to Equation (1), the welfare of f can then be written as:

$$\sum_{s \in \mathcal{S}} \max_{i=1}^n \sum_{j \in B_s} \widehat{v}_i(j).$$

This can be extended to arbitrary partitions of the items — i.e. partitions not necessarily respecting the bipartite graph — by letting $w_{sij} = \widehat{v}_i(j)$ for $(j, s) \in E$, and $w_{sij} = 0$ for $(j, s) \notin E$, and defining:

$$welfare(B) = \sum_{s \in \mathcal{S}} \max_{i=1}^n \sum_{j \in B_s} w_{sij}.$$

Since the weights w_{sij} do not reward assignments of items to signals not respecting the bipartite graph, finding the welfare-maximizing deterministic signaling scheme reduces to finding a partition B of items among signals maximizing $welfare(B)$.

⁷In *combinatorial auctions*, there is a set M of items and a set N of players. Each player is equipped with a *valuation* mapping subsets of M to the real numbers. Welfare maximization in combinatorial auctions is the problem of assigning the items to the players in order to maximize the sum of the players’ values for their assigned bundle of items.

⁸A set function $f : 2^X \rightarrow \mathbb{R}$ is called XOS if it can be written as the maximum of additive set functions. Specifically, if $f(A) = \max_{i=1}^k \sum_{j \in A} w_{ij}$ for some integer k , and weights $w_{ij} \in \mathbb{R}$ for $i \in [k]$ and $j \in X$. We say an XOS function is represented explicitly if the weights w_{ij} are given as input.

We observe that this is an instance of combinatorial auctions with XOS valuations. Namely, if we interpret the signals \mathcal{S} as the “bidders” in combinatorial auctions, and B_s as the “bundle” of items assigned to s , the valuation function of s is simply the XOS function $f(A) = \max_{i=1}^n \sum_{j \in A} w_{sij}$. Welfare maximization in combinatorial auctions, when players have XOS valuations written explicitly, admits an $e/(e-1)$ approximation algorithm that runs in polynomial time, as shown in [DS06]. When combined with Lemma 3.2, this proves Theorem 4.1 in the absence of a cardinality constraint.

4.1.2 The General Case: Reduction to Optimization over Mappings

We now consider the bipartite signaling problem with a cardinality constraint k on the number of signals used. As should be clear from Section 4.1.1, the general case of bipartite signaling reduces to a generalization of combinatorial auctions with XOS valuations — namely, with the additional constraint that at most k players win any items in the combinatorial auction. We are not aware of previous work on this problem, and therefore design an $e/(e-1)$ algorithm for bipartite signaling directly. We break our proof in two parts: first, in this section we show that computing an approximately welfare-maximizing scheme reduces to “guessing” the mapping from signals to winning players, then in Section 4.1.3 we show how to express optimization over these mappings as submodular function maximization subject to a *matroid*⁹ constraint, which admits an $e/(e-1)$ approximation algorithm by the result of [Von08].

We begin by observing that every signaling scheme induces a many-to-one mapping $w : \mathcal{S} \rightarrow [n] \cup \{*\}$, which maps each signal to the winning player given that signal, where $*$ denotes an unused signal. For deterministic signalling schemes that are valid for our instance of bipartite signalling, $w(s) = *$ for all but at most k signals s . We call such mappings w *feasible winner mappings*. We reduce the problem of computing a near optimal signalling scheme to an optimization problem over feasible winner mappings w .

Given a “guess” for the feasible mapping w associated with the welfare-maximizing signalling scheme, computing a deterministic optimal signalling scheme is trivial: each item $j \in \Omega$ is mapped to a winning player who likes it most subject to respecting the bipartite graph E . Formally, for every mapping $w : \mathcal{S} \rightarrow [n] \cup \{*\}$ where $w(s) \neq *$ for at most k signals, we let $f_w : \Omega \rightarrow \mathcal{S}$ be a deterministic signalling scheme satisfying

$$f_w(j) \in \operatorname{argmax}_{s:(j,s) \in E, w(s) \neq *} v_{w(s)}(j)$$

for items j where such a signal s exists, and $f_w(j)$ is the “no information” signal $s_0 \in \mathcal{S}$ otherwise (recall that s_0 is a valid signal for all items, see Section 2). The social welfare of signalling scheme f_w is, by appealing to Equation (1), at least

$$\text{welfare}(w) = \sum_{s:w(s) \neq *} \sum_{j \in f_w^{-1}(s)} \hat{v}_{w(s)}(j).$$

Reversing the order of summation, and observing that f_w maps j to a signal maximizing $v_{w(s)}(j)$ subject to respecting the bipartite graph, gives the following equivalent expression for $\text{welfare}(w)$:

⁹Recall a *matroid* (e.g. [Oxl92]) is a ground set X and a non-empty collection $\mathcal{I} \subseteq 2^X$ of *independent sets* such that: (i) whenever S is independent and $T \subseteq S$, T is also independent; (ii) whenever $S, T \in \mathcal{I}$ with $|T| < |S|$, there is some $x \in S \setminus T$ such that $T \cup \{x\} \in \mathcal{I}$.

$$welfare(w) = \sum_{j \in \Omega} \left(\max_{s: (j,s) \in E, w(s) \neq *} \widehat{v}_{w(s)}(j) \right). \quad (3)$$

The social welfare of f_w is at least $welfare(w)$. Note, however, that f_w may be infeasible, in that it may use up to $k + 1$ signals including s_0 — however, in that case $w(s_0) = *$, and therefore a feasible signalling scheme with welfare at least $welfare(w)$ can be gotten by simply choosing an arbitrary signal $s \neq s_0$ with $w(s) \neq *$, and reassigning all items $f_w^{-1}(s)$ to the “no information” signal s_0 . Moreover, when w is the mapping associated with the welfare-maximizing valid signalling scheme, f_w is a welfare maximizing valid signalling scheme with welfare exactly equal to $welfare(w)$. Therefore, finding an approximately welfare-maximizing signalling scheme reduces to finding a mapping w approximately maximizing $welfare(w)$. We summarize this in the following Lemma.

Lemma 4.2. *Fix $\alpha > 0$. Computing an α -approximately welfare maximizing bipartite signalling scheme reduces, in polynomial time, to computing a feasible winner mapping $w : \mathcal{S} \rightarrow [n] \cup \{*\}$ which α -approximately maximizes $welfare(w)$, as given in Equation (3).*

4.1.3 The General Case: Optimization over Feasible Winner Mappings

We now show how to find a feasible winner mapping $w : \mathcal{S} \rightarrow [n] \cup \{*\}$ which approximately maximizes $welfare(w)$, as given in Equation (3). We do so by posing this as a submodular function¹⁰ maximization problem subject to a *truncated partition matroid*¹¹ constraint, and invoking the result of Vondrak [Von08] which shows a polynomial-time $e/(e - 1)$ -approximation algorithm for submodular function maximization subject to an arbitrary matroid.

Recall that we defined a winner mapping w as feasible if $w(s) \neq *$ for at most k signals $s \in \mathcal{S}$. In order to pose our optimization problem as a constrained submodular maximization problem, we first relax the set of feasible mappings as follows. We consider many-to-many mappings $W \subseteq \mathcal{S} \times [n]$ from signals to players, and define the welfare of such a mapping W as follows:

$$welfare(W) = \sum_{j \in \Omega} \left(\max_{i, s: (j,s) \in E, (s,i) \in W} \widehat{v}_i(j) \right). \quad (4)$$

In other words, each item j may be assigned to a player i so long as there is a signal s that is valid for j , and i as one of the “winning” players of s as given by W ; the welfare-maximizing such assignment is used to calculate $welfare(W)$.

When W is a many-to-one mapping — i.e. assigns to each signal s at most one winning player $w(s)$ — it is easy to verify that $welfare(W)$ (Equation (4)) is equal to $welfare(w)$ (Equation (3)). Moreover, many-to-one maps W satisfying $|W| \leq k$ are in one-to-one correspondence with the set of feasible winner mappings $w : \mathcal{S} \rightarrow [n] \cup \{*\}$, where feasibility is as defined in Section 4.1.2. It is simple to verify that the family of subsets W of $\mathcal{S} \times [n]$ satisfying those two constraints is a *truncated partition matroid* on ground set $\mathcal{S} \times [n]$.

¹⁰A function $f : 2^U \rightarrow \mathbb{R}$ defined on all subsets of a finite non-empty set U is *submodular* if

$$f(S \cup \{i\}) - f(S) \leq f(T \cup \{i\}) - f(T)$$

for every $T \subseteq S \subset U$ and $i \notin S$.

¹¹A *truncated partition matroid* (see [Oxl92]) is a matroid (U, I) with the following structure. There exists a partition U_1, \dots, U_m of ground set U , integers k_1, \dots, k_m , and an integer k , such that a subset S of U is independent, i.e. is in I , if and only if $|S \cap U_j| \leq k_j$ for all $j \in 1, \dots, m$, and moreover $|S| \leq k$.

We now show that $welfare(W)$ is submodular. We define for each $j \in \Omega$ and $(s, i) \in [S] \times [n]$ a weight $w_{sij} = \widehat{v}_i(j)$ when $(j, s) \in E$, and $w_{sij} = 0$ otherwise. We then rewrite $welfare(W)$ as follows.

$$welfare(W) = \sum_{j \in \Omega} \max_{(s,i) \in W} w_{sij}. \quad (5)$$

Observe that $welfare(W)$ is the sum of $|\Omega|$ set functions, namely the functions $welfare_j(W) = \max_{(s,i) \in W} w_{sij}$ for all $j \in \Omega$. It is known, and easy to verify, that any set function $f : 2^X \rightarrow \mathbb{R}$ that simply associates a fixed weight w_x with each element $x \in X$, and evaluates to $f(S) = \max_{x \in S} w_x$, is submodular. Therefore, by extension, $welfare_j(W)$ is submodular for each j . It is also well known that submodular functions are closed under summation. Therefore, we conclude that $welfare(W)$ is submodular.

Since $welfare(W)$ is submodular, and moreover the set of W corresponding to feasible winner mappings is a matroid, invoking the the result of Vondrak [Von08] yields the followin Lemma.

Lemma 4.3. *There is a polynomial-time $e/(e-1)$ approximation algorithm for computing a feasible winner mapping w maximizing $welfare(w)$, as given in Equation (3).*

Combined with Lemma 4.2, this completes the proof of Theorem 4.1.

4.2 Revenue Maximization with Known Valuations

Next we show that, in the case of known valuations, there is a polynomial-time, constant factor approximation algorithm for revenue maximization in cardinality constrained signaling. Our algorithm (Algorithm 1) simply chooses the best of two signaling schemes, computed via procedures 2 and 3. We use OPTR and OPTW to denote the maximum revenue and welfare, respectively, of a cardinality-constrained signaling scheme, and $v_i(*) = \mathbb{E}_{j \sim p} v_i(j)$ to denote player i 's value for a random item. Moreover, we let $v^* = \max_i v_i(*)$ denote the maximum value of a player for a random item, and let i^* be the player attaining this maximum. We note that v^* upper-bounds the contribution of any individual player to the social welfare of any signaling scheme.

Input: Instance of cardinality constrained signaling, given by Ω, k , valuations v_1, \dots, v_n

Output: Signaling scheme x^*

- 1: Run Procedure 2 to compute deterministic signaling scheme g .
- 2: Run Procedure 3 to compute signaling scheme x .
- 3: Let x^* be the revenue-maximizing signaling scheme among g and x

Algorithm 1: Algorithm for Cardinality-constrained Signaling

We now provide some intuition for our algorithm. The signaling scheme computed by Procedure 2 guarantees near-optimal revenue when no individual player's contribution to the optimal social welfare is too large. An approximately welfare-maximizing signaling scheme f is computed as described in Section 4.1, then signals are sorted in decreasing order of their contribution to the social welfare, and then pairs of signals are "merged" in that order. Formally, merging two signals s and t in a deterministic signaling scheme $f : \Omega \rightarrow \mathcal{S}$ gives a new signaling scheme g with $g(j) = \{s, t\}$ whenever $f(j) = s$ or $f(j) = t$, and $g(j) = f(j)$ otherwise. Merging two signals s and t is tantamount to forcing the two winners of the auction in each of the signals to compete, extracting the value of at least one of them as revenue. Therefore, merging pairs of signals in order

Input: Instance of cardinality constrained signaling, given by Ω , k , valuations v_1, \dots, v_n

Output: Deterministic signaling scheme $g : \Omega \rightarrow [k]$

- 1: Compute a deterministic signaling scheme which approximately maximizes welfare (Section 4.1).
- 2: Repeatedly merge pairs of signals with the same winner, until each signal has a unique winner. Call the resulting signaling scheme f .
- 3: Sort the signals in decreasing order of their contribution to the social welfare of f : s_1, \dots, s_k
- 4: Let g be the signaling scheme which merges signals s_i and s_{i+1} in f , for all odd i .

Procedure 2: First Sub-procedure for Cardinality-constrained Signaling

Input: Instance of cardinality constrained signaling, given by Ω , k , and prior \mathcal{D} over valuations v_1, \dots, v_n

Output: Signaling scheme x

- 1: Exclude i^* , then compute a deterministic signaling scheme h approximately maximizing welfare for other players (Section 4.1). Let $welfare(h)$ denote its welfare (excluding i^*), and let $v_s = \max_{i \neq i^*} \sum_{j \in h^{-1}(s)} \widehat{v}_i(j)$ denote the contribution of signal s to the welfare.
- 2: Let $\alpha = v^*/welfare(h)$.
- 3: Let y be the signaling scheme which ignores the realization of the item, and outputs signal s with probability $\alpha \frac{v_s}{v^*}$. (It is easy to verify the probabilities sum to 1).
- 4: Let x be the signaling scheme which with probability 0.5 runs h , and with remaining probability 0.5 runs y .

Procedure 3: Second Sub-procedure for Cardinality-constrained Signaling

of their contribution to the social welfare extracts half the welfare of all but the most valuable signal in f . When no individual player contributes much to the social welfare, this is a constant-factor of the welfare of f .

The signaling scheme computed by Procedure 3, on the other hand, guarantees near-optimal revenue when a single player accounts for a large fraction of the optimal social welfare. In this case, i^* 's value for a random item, namely v^* , is on the order of the optimal social welfare. Procedure 3 first computes a deterministic signaling scheme h which $e/(e-1)$ -approximately maximizes welfare for players other than i^* . Then h is “mixed” with a signaling scheme y which releases no information — i.e. outputs a random signal independent of the realization of the item. Here, mixing two signaling scheme is defined as running each with equal probability. When the probability of each signal s in y is proportional to the contribution of s to the welfare of h , player i^* 's value conditioned on s is on the order of the value of the welfare of h conditioned on s . Therefore, i^* serves as a price-setting player for all signals, extracting a constant fraction of the social welfare of h in the process. Since, by Lemma 3.3, the welfare of h is a $(1 - 1/e)$ -approximation to the maximum revenue of any signaling scheme, this yields a constant approximation to the optimal revenue.

We now present a formal proof of the approximation ratio of our algorithm. The result follows from two Lemmas.

Lemma 4.4. *Fix $\beta > 0$. If $v^* \leq \beta OPTW$, then signaling scheme g — as computed by procedure 2 — has revenue at least $\frac{(1-1/e-\beta)}{2} OPTW$.*

Proof. Note that $welfare(f) \geq (1 - 1/e)OPTW$. It is an easy observation that Step 2, which merges signals with the same winning player, preserves the welfare of f .

Our assumption that any player's value for a random item is at most $\beta OPTW$ implies that signal s_1 accounts for no more than $\beta OPTW$ of the welfare of f — formally $\max_i \sum_{j \in f^{-1}(s)} \widehat{v}_i(j) \leq v^* \leq \beta OPTW$. Therefore, signals other than s_1 account for at least $(1 - 1/e - \beta)OPTW$ welfare in f .

Now, recall that g simply merges pairs of signals in f . Given two signals s and t in f , merged into a new signal $\{s, t\}$ in g , and an arbitrary player i , we observe that

$$\sum_{j \in g^{-1}(\{s, t\})} \widehat{v}_i(j) = \sum_{j \in f^{-1}(s)} \widehat{v}_i(j) + \sum_{j \in f^{-1}(t)} \widehat{v}_i(j).$$

Intuitively, the conditional value of player i for the new merged signal, weighted by the probability of the signal, is equal to the sum of i 's weighted conditional values for both the component signals. This implies that, given a merged signal $\{s, t\}$ in g , each of the (distinct) winners of s and t , which we denote by i_s and i_t , have at least as much (weighted) value for $\{s, t\}$ as they did for s and t individually. Since the player with the greatest weighted conditional value for a signal wins, and pays the second-greatest weighted value in expectation, we conclude that the contribution of $\{s, t\}$ to the revenue in g is at least the minimum of the contributions of s and t to the welfare of f .

Recalling that we sorted the signals in f in decreasing order of their contribution to the social welfare, and then merged them pairwise in that order, we conclude that the revenue of g is at least half the welfare of f after discarding s_1 . Our bound on the contribution of s_1 to the welfare of f then completes the proof. \square

Lemma 4.5. *If $v^* > \beta OPTW$, then x — as computed by procedure 3 — has revenue at least $\frac{\beta}{2}(1 - 1/e)OPTW$.*

of Lemma 4.5. First, observe that α , as stated in procedure 3, is at least β . Moreover, Lemma 3.3, combined with the result of Theorem 4.1, shows that $welfare(h) \geq (1 - 1/e)OPTW$. It remains to show that x extracts at least a $\alpha/2$ fraction of the welfare of h in revenue.

Let i_s denote the winning player of signal s in h (after excluding i^*). Recall that v_s , as defined in procedure 3, denotes the value of player i_s conditioned on signal s , weighted by the probability of the signal. Now consider signaling scheme x , which mixes — in equal measure — signaling scheme h with the scheme which outputs s with probability $\alpha \frac{v_s}{v^*}$. The (weighted) value of i_s for signal s in x is at least half what it was in h — namely $v_s/2$. Moreover, the (weighted) value of i^* for signal s is at least his value for a random item times the probability that signal s is output without regard to the realized item — doing the calculation, this is $v^* \cdot \frac{1}{2} \cdot \alpha \frac{v_s}{v^*} = \alpha v_s/2$.

Now, in x both i^* and i_s have (weighted) value for signal s equal to at least $\alpha v_s/2$. Therefore, the contribution of s to the revenue is at least this amount. Summing over all signals, the total revenue of x is at least $\frac{\alpha}{2} welfare(h)$. This completes the proof. \square

Setting $\beta = \frac{e-1}{2e-1}$ in Lemmas 4.4 and 4.5 proves the following theorem.

Theorem 4.6. *Algorithm 1 computes a $\frac{2e(2e-1)}{(e-1)^2} \approx 8.17$ approximation to the optimal revenue in the cardinality-constrained signaling problem.*

Note 4.7. *We note that optimizing Procedure 3 to mix h and y with probabilities $\gamma = \alpha/(1 + \alpha)$ and $(1 - \gamma)$, instead of 0.5 each, improves the approximation ratio in our theorem to 6.47. We omit the details.*

4.3 Unknown Valuations

We largely leave open the polynomial-time approximability of welfare and revenue maximization in the unknown valuations setting. Even in the case of welfare maximization in cardinality-constrained signaling, it appears challenging to design approximation algorithms with run-time scaling sub-exponentially in the support size of \mathcal{D} , the distribution over player valuation profiles. Nevertheless, we present here an extension of the result of Section 4.1.1 to distributions \mathcal{D} with constant size support. The runtime of the algorithm scales exponentially in the size t of the support. Moreover, we establish a connection of this problem to a generalization of the combinatorial auctions problem used in Section 4.1.1.

We assume the distribution \mathcal{D} over valuation profiles is listed explicitly as a set of matrices $\{v^1, \dots, v^t\}$ where $v^\ell \in \mathbb{R}^{[n] \times \Omega}$, and associated probabilities q_1, \dots, q_t . Using a similar derivation to that in Section 4.1, the welfare of a deterministic signaling scheme f can be written as

$$\sum_{s \in \mathcal{S}} \sum_{\ell=1}^t q_\ell \max_{i=1}^n \sum_{j \in f^{-1}(s)} \widehat{v}_i^\ell(j).$$

Recall that, as described in Section 4.1.1, the variant with $t = 1$ can be interpreted as an instance of combinatorial auctions with players corresponding to \mathcal{S} , who are equipped with XOS valuations. Similarly, the variant with general t can be interpreted as an instance of combinatorial auctions where each player's valuation is a sum of XOS functions. We are not aware of any non-trivial polynomial-time algorithms for this problem that run in time polynomial in t in the computational complexity model. However, since XOS valuations, and therefore their sums, are *subadditive*, the algorithm of Feige [Fei06] guarantees a 2-approximation if players can answer *demand queries*. Unfortunately we can show that, unlike for XOS functions, answering a demand query for a sum of XOS functions, even approximately to within a constant factor, is NP-hard — the proof is deferred to the full version of this paper. However, this does not rule out other approaches to our problem not involving demand queries.

We now show that a $1 - 1/e$ approximation is possible in time exponential in t , and polynomial in all other parameters of an instance. We rewrite the welfare of a signaling scheme f as follows.

$$\sum_{s \in \mathcal{S}} \max_{i_1, i_2, \dots, i_t} \sum_{j \in f^{-1}(s)} \sum_{\ell=1}^t q_\ell \widehat{v}_{i_\ell}^\ell(j).$$

This, again, is an instance of combinatorial auctions with XOS valuations, though each XOS function is the maximum of n^t additive functions. The $e/(e - 1)$ -approximation algorithm of [DS06] then runs in time polynomial in n^t .

4.4 Hardness of Approximation

We now show that our approximation ratio for welfare-maximization in bipartite signaling is tight, even for the special case of cardinality-constrained signaling and known valuations. As a corollary, the same hardness of approximation result holds for revenue maximization. We use an approximation-preserving reduction from the APX-hard problem max-cover. An instance of max-cover is given by a ground set $[m]$, a family A_1, \dots, A_n of subsets of $[m]$, and an integer k . The goal is to find k sets from A_1, \dots, A_n whose union is as large as possible.

Given an instance of max cover, we construct an instance of cardinality-constrained signaling with known valuations as follows. We let the set Ω of items be $[m]$, and associate with each set A_i a player i with valuation $v_i : [m] \rightarrow \{0, 1\}$ such that $v_i(j) = 1$ if and only if $j \in A_i$. Moreover, we let k be the constraint on the number of allowable signals, and let p be the uniform distribution over Ω .

Given a solution A_{i_1}, \dots, A_{i_k} for max-cover covering m' items, we show a signaling scheme with welfare at least m'/m . We assign items in A_{i_1} to signal 1, then assign items in $A_{i_2} \setminus A_{i_1}$ to signal 2, and so on until signal k . We observe that, for each signal $\ell \in [k]$, there is at least one player — in particular player i_ℓ — who values at least $|A_{i_\ell} \setminus \cup_{\ell' < \ell} A_{i_{\ell'}}|$ of the items assigned to signal ℓ . This implies that the welfare of the signaling scheme is at least $\frac{1}{m} |\cup_{\ell=1}^k A_{i_\ell}| = m'/m$.

Conversely, given a signaling scheme with welfare m'/m , we show a solution to max-cover with coverage of at least m' . Each signaling scheme partitions the items into k subsets B_1, \dots, B_k , each of which — say B_ℓ — is associated with a winning player i_ℓ . The welfare of the signaling scheme is the sum, over all winning players i_ℓ , of the number of items in B_ℓ valued by i_ℓ , scaled by the per-item probability of $\frac{1}{m}$. Formally, the welfare is $\frac{1}{m} \sum_{\ell=1}^k |A_{i_\ell} \cap B_\ell|$. Letting $m' = \sum_{\ell=1}^k |A_{i_\ell} \cap B_\ell|$, it is clear from the fact that B_1, \dots, B_k is a partition of the items that $|\cup_{\ell=1}^k A_{i_\ell}| \geq m'$, as needed to show a solution to max-cover with coverage at least m' .

The above reduction, coupled with the hardness of approximation result for max-cover given in [Fei98], yields the following theorem.

Theorem 4.8. *There is no polynomial-time c -approximation algorithm for welfare-maximization in cardinality-constrained signaling with known valuations, for any constant $c < \frac{e}{e-1}$, unless $P = NP$.*

The hardness result of Theorem 4.8 also holds for revenue maximization. This follows from a simple reduction from welfare maximization to revenue maximization: given an instance of cardinality constrained signaling, produce a new instance whose welfare and revenue are equal by simply duplicating each player.

Corollary 4.9. *There is no polynomial-time c -approximation algorithm for revenue-maximization in cardinality-constrained signaling with known valuations, for any constant $c < \frac{e}{e-1}$, unless $P = NP$.*

5 Communication Bounded Signaling Schemes

In this section, we consider signaling schemes in which items $y \in \Omega$ correspond to points in d dimensional Euclidean space \mathbb{R}^d . In such settings, we consider valuation functions which are also parameterized by points $v_i \in \mathbb{R}^d$. Natural valuation functions then include inner products (i.e. $v_i(y) = \langle y, v_i \rangle$) and distances (i.e. $v_i(y) = \|y - v_i\|$). We think of d as being very large, and so we will be concerned with *space bounded signaling schemes*, defined to be schemes that can communicate only a bounded number of bits per signal.

Definition 5.1. *A finite set of signals S has b -bounded length if $\log |S| \leq b$. The set of b -bounded length signaling schemes is $\mathcal{F} \subseteq S^\Omega$ such that for each $f \in \mathcal{F}$: the set of all possible signals generated by f (i.e. $\cup_{y \in \Omega} f(y)$) has b bounded length. We call such an $f \in \mathcal{F}$ a b -bounded length signaling scheme.*

Remark 5.2. *Note that any signal in a b -bounded length signaling scheme can be transmitted using at most b bits.*

5.1 Inner Product Valuations

Let $\Omega = \{y \in \mathbb{R}_+^d : \|y\|_1 = 1\}$ be the set of d dimensional non-negative real vectors with ℓ_1 weight 1. These can be thought of as distributions over d “features” which describe the product y . Individuals i have valuation functions $v_i : \Omega \rightarrow \mathbb{R}$ parameterized by (abusing notation) a vector $v_i \in \mathbb{R}^d$ with $\|v_i\|_\infty \leq 1$. Bidder i ’s valuation for good y is defined to be $v_i(y) \equiv \langle v_i, y \rangle$. Intuitively, each component j of v_i represents agent i ’s affinity for feature j , which can range in $[-1, 1]$. $v_i(y)$ is agent i ’s average affinity for the distribution over features represented by good y . We wish to design good b -bounded length signaling schemes for inner product valuations, for b as small as possible.

5.1.1 Multiplicative Weights Signaling

In this section, we use the versatile multiplicative weights framework of [AHK12] to give a communication bounded signaling scheme. The idea is the following: For simplicity, we first consider the known valuation setting in which the signaling scheme can be parameterized by the actual valuations v_1, \dots, v_n of the bidders. When a good $y \in \Omega$ arrives, we use multiplicative weights to learn an ϵ -approximate representation \hat{y} with respect to the bidder valuation functions v_1, \dots, v_n – that is, a representation \hat{y} such that for all i , $|v_i(y) - v_i(\hat{y})| \leq \epsilon$. We can do this by updating multiplicative weights at most $O(\log d/\epsilon^2)$ times using the valuation functions themselves as loss functions. \hat{y} would itself therefore make a terrific signal – it would approximately represent every bidder’s valuation for the good y . However, \hat{y} is also a vector in \mathbb{R}^d , and so it is not clear why we should be able to communicate it in a space bounded way. The key insight is that it is not necessary to communicate \hat{y} directly, but merely communicate which collection of valuation functions were used to update multiplicative weights when learning \hat{y} – using this information, each bidder can reconstruct \hat{y} for themselves. (Of course, this “reconstruction” can be automated, so the bidders can still see a natural signal). Since each valuation function can be indexed with only $O(\log n)$ bits, and there are only $O(\log d/\epsilon^2)$ updates in total, this gives a $O(\log n \log d/\epsilon^2)$ -space bounded signaling scheme that approximates the optimal welfare within an additive loss of ϵ . We then extend this to the case in which the valuation functions v_i are not known, but instead drawn i.i.d. from a known prior \mathcal{P} . This extension involves parameterizing the same multiplicative weights signaling scheme with m i.i.d. samples from \mathcal{P} . m here corresponds to the *sample complexity* of the corresponding learning problem on linear valuations¹² and the space used by this signaling scheme depends on $\log m$ (since we must index updates from this set of m vectors). We now define our signaling scheme formally.

We define a signaling scheme $f_{\text{MW}, \epsilon, z_1, \dots, z_m}$ parameterized by a no regret algorithm (in this case the multiplicative weights algorithm MW), an accuracy parameter ϵ , and m vectors $z_1, \dots, z_m \in \mathbb{R}^d$. The parameters of the signaling scheme will be public knowledge, and it will be used to generate b -bounded length signals as follows:

Theorem 5.3. *For any ϵ , any vectors z_1, \dots, z_m such that for all i , $\|z_i\|_\infty \leq 1$, and any $y \in \Omega$, $f_{\text{MW}, \epsilon, z_1, \dots, z_m}$ runs for $T \leq \frac{16 \log d}{\epsilon^2}$ rounds.*

Proof. $f_{\text{MW}, \epsilon, z_1, \dots, z_m}(y)$ runs an instantiation of the Multiplicative Weights Framework for Arora, Hazan, and Kale [AHK12] for T rounds, with update parameter $\epsilon/4$, and loss vectors at each round

¹²The *sample complexity* of a learning problem is, informally, the number of samples that need to be drawn from a distribution \mathcal{P} , such that if we learn a hypothesis that is consistent on the sampled points, then with high probability, the hypothesis is consistent on *new* points drawn from the same distribution.

Input: An instance $y \in \Omega$.

Output: A bounded length signal s .

- 1: Initialize $\hat{y}^1 \in \mathbb{R}^d$ such that $\hat{y}_j^1 = 1/d$ for all $j \in [d]$.
- 2: Initialize $T \leftarrow 1$.
- 3: **while** there exists an index i such that $|\langle z_i, y \rangle - \langle z_i, \hat{y}^T \rangle| \geq \epsilon/2$ **do**
- 4: Let $\text{Update}_T \leftarrow i$, $\text{Sign}_T \leftarrow \text{sign}(\langle z_i, \hat{y}^T \rangle - \langle z_i, y \rangle)$,
- 5: For all j let $\hat{y}_j^{T+1} \leftarrow \hat{y}_j^T \cdot (1 - \text{Sign}_T \cdot \frac{\epsilon}{4} \cdot z_{i,j})$.
- 6: Normalize \hat{y}^{T+1} such that $\|\hat{y}^{T+1}\|_1 = 1$. Let $T \leftarrow T + 1$
- 7: **end while**
- 8: Let $s = ((\text{Update}_1, \text{Sign}_1), \dots, (\text{Update}_T, \text{Sign}_T))$

Algorithm 4: Algorithm for computing the signal $f_{\text{MW}, \epsilon, z_1, \dots, z_m}$

t defined to be $\ell^t = \text{Sign}_t \cdot z_{\text{Update}_t}$. By the regret bound of Multiplicative Weights (see e.g. [AHK12] Corollary 2.2), we have for all $x \in \mathbb{R}_+^d$ with $\|x\|_1 = 1$:

$$\sum_{t=1}^T (\langle \ell^t, \hat{y}^t \rangle - \langle \ell^t, x \rangle) \leq \frac{\epsilon}{4} \sum_{t=1}^T |\langle \ell^t, x \rangle| + \frac{4 \ln d}{\epsilon}$$

where $|\ell^t|$ denotes a coordinate-wise absolute value. Note that this corresponds to the standard “no regret bound” that readers may be more familiar with: here the “experts” correspond to the d standard basis vectors, x corresponds to a distribution over experts, and the above bound simply states that multiplicative weights achieves diminishing regret with respect to the best expert (and therefore with respect to any distribution over experts). Taking $x = y$, we note that the loss vectors $\ell^t = \text{Sign}_t \cdot z_{\text{Update}_t}$ have been constructed such that at every round t , $\langle \ell^t, \hat{y}^t \rangle - \langle \ell^t, y \rangle \geq \epsilon/2$. Moreover, since for all i $\|z_i\|_\infty \leq 1$, by definition for each coordinate j , $|\ell_j^t| \leq 1$. Therefore, for each t , $\langle |\ell^t|, y \rangle \leq 1$ since $\|y\|_1 = 1$. Therefore, the above bound becomes:

$$\frac{T\epsilon}{2} \leq \frac{T\epsilon}{4} + \frac{4 \ln d}{\epsilon}$$

Solving for T , we find that it must be that $T \leq \frac{16 \log d}{\epsilon^2}$ as desired. \square

We now make two observations. The first is that $f_{\text{MW}, \epsilon, z_1, \dots, z_m}(y)$ produces bounded length signals:

Corollary 5.4. *For any ϵ , any vectors z_1, \dots, z_m such that for all i , $\|z_i\|_\infty \leq 1$, and any $y \in \Omega$, $f_{\text{MW}, \epsilon, z_1, \dots, z_m}$ is a b -bounded length signaling scheme for $b = \left(\frac{16 \log d (\log m + 1)}{\epsilon^2} \right)$.*

Proof. For $t = 1, \dots, T$, Update_t can be communicated with $\log m$ bits and Sign_t can be communicated by 1 bit. By Theorem 5.3, for all $y \in \Omega$, $T \leq \frac{16 \log d}{\epsilon^2}$ \square

The next is that the signal $s = ((\text{Update}_1, \text{Sign}_1), \dots, (\text{Update}_T, \text{Sign}_T))$ is sufficient for each agent i to reconstruct \hat{y}^{T+1} :

Observation 5.5. $\hat{y}^{T+1} \equiv \hat{y}^{T+1}(s)$ is a function only of the signal $s = ((\text{Update}_1, \text{Sign}_1), \dots, (\text{Update}_T, \text{Sign}_T))$.

Given this observation, it is helpful to think about the signal “really being” the vector \hat{y}^{T+1} that well approximates $\langle y, z_i \rangle$ for all i . The signal s is just a concise way of transmitting this vector.

First, we show that when we have known valuations, the multiplicative weights signaling scheme is competitive with the optimal b -bounded length signaling scheme.

Theorem 5.6. *Let $OPT_U = \sum_{y \in \Omega} \Pr_{j \sim p}[j = y] \cdot \max_{i \in [n]} v_i(y)$ denote the optimal social welfare in the unconstrained setting. In the known valuation setting, the welfare obtained by the multiplicative weights signaling scheme given vectors $(z_1, \dots, z_n) \equiv (v_1, \dots, v_n)$ is:*

$$\text{welfare}(f_{MW, \epsilon, v_1, \dots, v_n}, v) \geq OPT_U - \epsilon$$

In particular, it is within ϵ of the optimal welfare obtained by any bounded length signaling scheme.

Proof. Fix any $y \in \Omega$, consider $s = s(y) = f_{MW, \epsilon, v_1, \dots, v_n}(y)$, and let $\hat{y}^{T+1} \equiv \hat{y}^{T+1}(s)$. For any bidder i , let $S = \{y : |\langle y, v_i \rangle - \langle \hat{y}^{T+1}(s), v_i \rangle| \leq \epsilon/2\}$ be the set of goods whose value to bidder i differs by less than ϵ from the value of good \hat{y}^{T+1} . Note that by the construction of \hat{y}^{T+1} , it must be the case that $\Pr_{j \sim p}[j \notin S | s] = 0$ (because otherwise the multiplicative weights signaling scheme would not have halted). Therefore, we have $\Pr_{j \sim p}[j \in S | s] = 1$ and we can calculate:

$$\begin{aligned} v_i | s(y) &= \mathbb{E}_{j \sim p}[v_i(j) | f(j) = s] = \sum_{j \in \Omega} v_i(j) \cdot \Pr_p[j | f(j) = s] \\ &\geq \sum_{j \in S} v_i(j) \cdot \Pr_p[j | f(j) = s] \geq \min_{j \in S} v_i(j) \cdot \Pr_p[j \in S | f(j) = s] \geq (v_i(y) - \epsilon) \cdot 1 \end{aligned}$$

Finally, we can lower bound the expected welfare of the multiplicative weights signaling scheme, and compare it to the unconstrained optimal welfare. For all v :

$$\begin{aligned} \text{welfare}(f_{MW, \epsilon, v_1, \dots, v_n}, v) &= \sum_{y \in \Omega} \Pr_{j \sim p}[j = y] \cdot \max_{i \in [n]} v_i | s(y) \\ &\geq \sum_{y \in \Omega} \Pr_{j \sim p}[j = y] \cdot \max_{i \in [n]} (v_i(y) - \epsilon) = OPT_U - \epsilon \end{aligned}$$

□

We now adapt our signaling scheme slightly, and show that it works not just in the known valuation setting, but also in the Bayesian setting when there is a prior \mathcal{P} from which the valuations v_i of the agents are drawn i.i.d. Consider the slightly modified family of signaling schemes f' , still defined in terms of an accuracy parameter ϵ and m vectors z_1, \dots, z_m . The difference between f' and f is only that f' updates the multiplicative weights hypothesis at most 1 time for every vector z_i . This difference is not required for the accuracy theorem that we prove (i.e. we could just as well have used the signaling scheme defined by f), but simplifies the analysis.

Because we continue to have that \hat{y}^T is updated only on vectors z_i such that $|\langle z_i, y \rangle - \langle z_i, \hat{y}^T \rangle| \geq \epsilon/2$, our bound on the number of updates is identical as it was for signaling scheme f , and so we have an identical corollary:

Corollary 5.7. *For any ϵ , any vectors z_1, \dots, z_m such that for all i , $\|z_i\|_\infty \leq 1$, and any $y \in \Omega$, $f'_{MW, \epsilon, z_1, \dots, z_m}$ is a b -bounded length signaling scheme for $b = \left(\frac{16 \log d \log m + 1}{\epsilon^2}\right)$.*

Input: An instance $y \in \Omega$.

Output: A bounded length signal s .

- 1: Initialize $\hat{y}^1 \in \mathbb{R}^d$ such that $\hat{y}_j^1 = 1/d$ for all $j \in [d]$.
- 2: Initialize $T \leftarrow 1$.
- 3: **for** $i = 1$ to m **do**
- 4: **if** $|\langle z_i, y \rangle - \langle z_i, \hat{y}^T \rangle| \geq \epsilon/2$ **then**
- 5: Let $\text{Update}_T \leftarrow i$, $\text{Sign}_T \leftarrow \text{sign}(\langle z_i, \hat{y}^T \rangle - \langle z_i, y \rangle)$,
- 6: For all j let $\hat{y}_j^{T+1} \leftarrow \hat{y}_j^T \cdot (1 - \text{Sign}_T \cdot \frac{\epsilon}{2} \cdot z_{i,j})$.
- 7: Normalize \hat{y}^{T+1} such that $\|\hat{y}^{T+1}\|_1 = 1$. Let $T \leftarrow T + 1$
- 8: **end if**
- 9: **end for**
- 10: Let $s = ((\text{Update}_1, \text{Sign}_1), \dots, (\text{Update}_T, \text{Sign}_T))$

Algorithm 5: Algorithm for computing the signal $f'_{\text{MW}, \epsilon, z_1, \dots, z_m}$

We now argue that if f is parameterized with m vectors z_i drawn i.i.d. from \mathcal{P} (for sufficiently large m), then in fact with high probability, $f'_{\text{MW}, \epsilon, z_1, \dots, z_m}$ is a competitive signaling scheme for the actual agent valuations v_1, \dots, v_n , whenever v_1, \dots, v_n are also drawn i.i.d. from \mathcal{P} .

Theorem 5.8. *Let*

$$m = \frac{2n \left(\frac{16 \log d}{\epsilon^2} + 1 \right) \left(\log \left(\frac{16 \log d}{\epsilon^2} \right) + \log \frac{2n}{\delta} \right)}{\delta} = \tilde{O} \left(\frac{n \log d}{\delta \epsilon^2} \right).$$

Fix any $y \in \Omega$, and let z_1, \dots, z_m and v_1, \dots, v_n be i.i.d. draws from \mathcal{P} . Let $s = f'_{\text{MW}, \epsilon, z_1, \dots, z_m}(y)$, and define $S(v, s) = \{y : |\langle y, v \rangle - \langle \hat{y}^{T+1}(s), v \rangle| \leq \epsilon/2\}$. Then we have:

$$\Pr_{z_1, \dots, z_m, v \sim \mathcal{D}} [y \notin S(v, s)] \leq \frac{\delta}{n}$$

Proof. We will show that for any $y \in \Omega$, with probability $1 - \delta/2n$ over the choice of z_1, \dots, z_m , s is such that: $\Pr_{v \sim \mathcal{P}} [y \notin S(v, s)] \leq \delta/2n$, which is enough to prove our claim. Let $U = \frac{16 \log d}{\epsilon^2}$ be the maximum number of update rounds that $f'_{\text{MW}, \epsilon, z_1, \dots, z_m}$ can ever conduct, as bounded by Theorem 5.3. Since $f'_{\text{MW}, \epsilon, z_1, \dots, z_m}$ considers running an update on each z_i in sequence but never conducts more than U updates, there must be some consecutive sequence z_i, \dots, z_j of length at least $(j-i+1) \geq \frac{m}{U}$ on which no updates are performed. We will show that if $\Pr_{v \sim \mathcal{D}} [y \notin S(v, s)] \geq \delta/2n$, then the probability of this event occurring is at most $\delta/2n$. Note that if $\Pr_{v \sim \mathcal{D}} [y \notin S(v, s)] \geq \delta/2n$, since each z_i is independently sampled from \mathcal{P} , the probability of an update occurring at round i is at least $\delta/2n$. Therefore, the probability of there existing such a long sequence between updates is at most:

$$(U+1) \left(1 - \frac{\delta}{2n}\right)^{m/U} \leq (U+1) e^{-m\delta/2nU}$$

Setting m as in the theorem statement makes this probability at most $\delta/2n$ as desired. \square

A simple corollary of Theorem 5.8 is an analogous welfare guarantee for the multiplicative weights signaling scheme in the Bayesian setting:

Theorem 5.9. Let $OPT_U = \sum_{y \in \Omega} \Pr_{j \sim p}[j = y] \cdot \max_{i \in [n]} v_i(y)$ denote the optimal social welfare in the unconstrained setting with known valuations. In the Bayesian setting in which each $v_i \sim \mathcal{P}$ is sampled i.i.d. from a known prior \mathcal{P} , the welfare obtained by the multiplicative weights signaling scheme given vectors (z_1, \dots, z_m) sampled independently from the prior distribution \mathcal{P} is:

$$\text{welfare}(f'_{MW, \epsilon, v_1, \dots, v_m}, v) \geq OPT_U - \epsilon - \delta$$

where $m = \tilde{O}\left(\frac{n \log d}{\delta \epsilon^2}\right)$. In particular, it is a $16 \log d (\log m + 1) / \epsilon^2$ bounded length signaling scheme that obtains welfare within $\epsilon + \delta$ of the optimal welfare obtained by any bounded length signaling scheme.

Proof. Independently for every $y \in \Omega$, consider $s = s(y) = f_{MW, \epsilon, v_1, \dots, v_m}(y)$, and let $\hat{y}^{T+1} \equiv \hat{y}^{T+1}(s)$. For any bidder i , let $S_i = \{y : |\langle y, v_i \rangle - \langle \hat{y}^{T+1}(s), v_i \rangle| \leq \epsilon/2\}$ be the set of goods whose value to bidder i differs by less than ϵ from the value of good \hat{y}^{T+1} . By Theorem 5.8, $\Pr_{j \sim p}[j \in S_i | s] \geq 1 - \delta/n$ and so by a union bound $\Pr_{j \sim p}[\forall i, j \in S_i | s] \geq 1 - \delta$ and we can calculate for all i :

$$\begin{aligned} v_i | s(y) &= \mathbb{E}_{j \sim p}[v_i(j) | f(j) = s] \\ &= \sum_{j \in \Omega} v_i(j) \cdot \Pr_p[j | f(j) = s] \\ &\geq \sum_{j \in S} v_i(j) \cdot \Pr_p[j | f(j) = s] \\ &\geq \min_{j \in S} v_i(j) \cdot \Pr_p[j \in S | f(j) = s] \\ &\geq (v_i(y) - \epsilon) \cdot (1 - \delta) \\ &\geq v_i(y) - \epsilon - \delta \end{aligned}$$

where the last inequality follows from the fact that $v_i(y) \leq 1$. Finally, we can lower bound the expected revenue of the multiplicative weights signaling scheme, and compare it to the unconstrained optimal welfare. For all v :

$$\begin{aligned} \text{welfare}(f'_{MW, \epsilon, z_1, \dots, z_m}, v) &= \sum_{y \in \Omega} \Pr_{j \sim p}[j = y] \cdot \max_{i \in [n]} v_i | s(y) \\ &\geq \sum_{y \in \Omega} \Pr_{j \sim p}[j = y] \cdot \max_{i \in [n]} (v_i(y) - \epsilon - \delta) \\ &= OPT_U - \epsilon - \delta \end{aligned}$$

□

5.2 Subspace Valuations

In this section, we use a version of the Johnson Lindenstrauss lemma, for matrices implicitly defined by limited independence families of hash functions, to give bounded space signaling schemes for “subspace valuations”. Subspace valuations are defined over a point set Ω of unit vectors in Euclidean space, and can be seen as a generalization of the “inner product valuations” considered

in the previous section¹³. An agent may specify up to k points in Ω , which indicates that he is equally happy with any linear combination of these k points (i.e. his k points define a subspace). His value for a good is defined to be its distance to this subspace. The Johnson Lindenstrauss lemma lets us take a “projection” of the point into a lower dimensional space, in such a way that with high probability, each agent is able to estimate its value with high probability. We use the fact that limited independence JL matrices can be concisely represented – this allows us to construct a new projection matrix for every good, and include its description as part of our signal. This allows us to give a strong Bayesian guarantee – our algorithm achieves close to optimal welfare against any (arbitrarily correlated) prior value distribution, even without knowledge of the distribution. Indeed, our algorithm works even against *adversarially selected* goods $y \in \Omega$, that need not be drawn from any distribution! We state our main theorem here but leave the details to the appendix.

Let $\Omega = \{y \in \mathbb{R}^d : \|y\|_2 = 1\}$ be the set of d dimensional unit vectors in Euclidean space. Individuals i have valuation functions parameterized by ℓ_i orthogonal unit vectors $z_1^i, \dots, z_{\ell_i}^i$ for some $\ell_i \leq k$. These vectors define a subspace $S_i \equiv \text{span}(z_1^i, \dots, z_{\ell_i}^i)$, and the value that agent i has for a good y is the distance between y and agent i 's subspace S_i :

$$v_i(y) \equiv 1 - d(y, S_i) = 1 - \min_{x \in S_i} \|x - y\|_2$$

For each agent i , we can think of the vectors $z_1^i, \dots, z_{\ell_i}^i$ as specifying up to k “ideal” goods, and that agent i is equally happy with any linear combination of his ideals. His valuation for a good y drops off with the distance from y to his set of ideal goods. Note that it is not necessary that agent i actually specify *orthogonal goods* – if his ideal goods $z_1^i, \dots, z_{\ell_i}^i$ are not orthogonal, we can simply orthonormalize them (using, say, the Gram Schmidt algorithm), which does not alter the subspace that they define.

We recall that for any y , $\arg \min_{x \in S_i} \|x - y\|_2 = \sum_{j=1}^{\ell_i} \langle y, z_j^i \rangle \cdot z_j$, and so we can write: $d(y, S_i) = \sqrt{1 - \sum_{j=1}^{\ell_i} \langle x, z_j^i \rangle^2}$.

To achieve a low space signaling scheme for subspace valuations, we make use of the Johnson-Lindenstrauss lemma.

We will use (a corollary of) a limited-independence version of the Johnson-Lindenstrauss lemma presented in [KN10], first proven by [Ach01, CW09]. This version of the lemma holds even for concisely represented projection matrices, which allows us to communicate the projection matrix itself as part of our signal. The advantage of doing this is that our algorithm will get strong utility guarantees even in the *prior free* setting, when valuation functions can be drawn from a worst-case (arbitrarily correlated) prior on distributions, and goods can be selected by an adversary, rather than being drawn from any distribution.

Corollary 5.10 (The Johnson-Lindenstrauss Lemma with Limited Independence [Ach01, CW09, KN10]). *For $d > 0$ an integer and any $0 < \epsilon, \delta < 1/2$, let A be a $T \times d$ random matrix with $\pm 1/\sqrt{T}$ entries that are r -wise independent for $T \geq 4 \cdot 64^2 \epsilon^{-2} \log(1/\delta)$ and $r \geq 2 \log(1/\delta)$. Then for any $x, y \in \mathbb{R}^d$:*

$$\Pr_A [|\langle (Ax), (Ay) \rangle - \langle x, y \rangle| \geq \frac{\epsilon}{2} (\|x\|_2^2 + \|y\|_2^2)] \leq 2\delta$$

¹³The class of valuations may be seen as a generalization of inner product valuations, but the signaling scheme here works in a different range of parameters as the multiplicative weights signaling scheme. Specifically, for the multiplicative weights signaling scheme, we assumed that Ω consisted of the set of unit vectors in ℓ_1 space, and that valuation functions were defined by unit vectors in ℓ_∞ space. In this section, both points $y \in \Omega$, and the vectors which parameterize agent valuation functions are unit vectors in ℓ_2 space.

Input: An instance $y \in \Omega$.

Output: A bounded length signal s .

- 1: Generate a $T \times d$ random $\pm 1/\sqrt{T}$ -valued matrix A with r -wise independent entries for $r = 2 \log(3nk/\epsilon)$ and $T = \frac{131072k^2 \log(3n/\epsilon)}{\epsilon^4}$.
- 2: Let $\hat{y} = Ay$ and let $\hat{y}' = \hat{y}$ discretized to $\log(3d/\epsilon)$ bits of precision.
- 3: Let $s = (A, \hat{y}')$.

Algorithm 6: Algorithm for computing the signal $f_{\text{JL},k,\epsilon}$

Remark 5.11. *The matrix A in our algorithm will be implicitly represented by a hash function mapping coordinates of the matrix A to their values. There are various ways to select a hash function from a family of r -wise independent hash functions mapping $[T \times d] \rightarrow \{0, 1\}$. The simplest, and one that suffices for our purposes, is to select the smallest integer s such that $2^s \geq T \times d$, and then to let g be a random degree r polynomial in the finite field $\mathbb{GF}[2^s]$. Selecting and representing such a function takes time and space $O(r \cdot s) = O(r(\log d + \log T))$. g is then an unbiased r -wise independent hash function mapping $\mathbb{GF}[2^s] \rightarrow \mathbb{GF}[2^s]$. Taking only the last output bit gives an unbiased r -wise independent hash function mapping $[s \times d]$ to $\{0, 1\}$, as desired.*

We first observe that $f_{\text{JL},k,\epsilon}$ generates bounded length signals.

Observation 5.12. *$f_{\text{JL},k,\epsilon}$ is an $O\left(\frac{k^2 \log(n/\epsilon) \log(d/\epsilon)}{\epsilon^4} (\log d + \log(k \log(n/\epsilon)))\right)$ -bounded length signaling scheme.*

Proof. This follows directly from our choice of r and T , the fact that A can be represented using $O(r(\log d + \log T))$ bits, and the fact that we discretize each coordinate of \hat{y}' to $\log(3d/\epsilon)$ bits of precision. \square

We now show that $f_{\text{JL},k,\epsilon}$ is welfare competitive with the optimal unconstrained signaling scheme.

Theorem 5.13. *Let $OPT_U = \sum_{y \in \Omega} \Pr_{j \sim p}[j = y] \cdot \max_{i \in [n]} v_i(y)$ denote the optimal social welfare in the unconstrained setting. For every distribution \mathcal{D} over subspace valuation functions and every distribution p over goods $y \in \Omega$,*

$$\text{welfare}(f_{\text{JL},k,\epsilon}, v) \geq OPT_U - \epsilon$$

In fact, for every v , this guarantee holds pointwise for goods $y \in \Omega$ even if adversarially chosen. For all $y \in \Omega$, for $s = f_{\text{JL},k,\epsilon}(y)$:

$$\max_{i \in [n]} \sum_{j \in \Omega} x(j, s) p_j v_i(j) \geq \max_{i \in [n]} v_i(y) - \epsilon$$

Proof. It suffices to prove the second, stronger claim. Fix a good $y \in \Omega$ and let $f_{\text{JL},k,\epsilon}(y) = s \equiv (A, \hat{y}')$. For each bidder i , define $S_i(A, \hat{y}, \delta) = \{x \in \Omega : |\sqrt{1 - \sum_{j=1}^{\ell_i} \langle x, z_j^i \rangle^2} - \sqrt{1 - \sum_{j=1}^{\ell_i} \langle \hat{y}, Az_j^i \rangle^2}| \leq \delta\}$. By our choice of T , we have: $\Pr[y \notin S_i(A, \hat{y}, \epsilon/3)] \leq \epsilon/3n$, where the probability is taken over the choice of projection matrix A . Moreover, by our choice of discretization, we have $S_i(A, \hat{y}, \epsilon/3) \subset$

$S_i(A, \hat{y}', 2\epsilon/3)$. Therefore by a union bound: $\Pr[\exists i : y \notin S_i(A, \hat{y}', 2\epsilon/3)] \leq \epsilon/3$. Therefore, for all i :

$$\begin{aligned} \sum_{j \in \Omega} x(j, s) p_j v_i(j) &\geq \sum_{j \in S_i(A, \hat{y}', 2\epsilon/3)} x(j, s) p_j v_i(j) \\ &\geq \min_{j \in S_i(A, \hat{y}', 2\epsilon/3)} v_i(j) \sum_{j \in S_i(A, \hat{y}', 2\epsilon/3)} x(j, s) p_j \\ &\geq (v_i(y) - 2\epsilon/3) \Pr[y \in S_i(A, \hat{y}', 2\epsilon/3)] \\ &\geq v_i(y) - \epsilon \end{aligned}$$

□

6 Conclusions and Future Work

Our results initiate the study of signaling in constrained settings. Whereas we obtain preliminary positive and negative results for some natural problems in this setting, we leave open a rich selection of algorithmic problems. We conclude with a statement of several of these open questions.

We leave several open questions in the bipartite signaling setting. Is there a constant-factor approximation algorithm for revenue maximization in this setting with known valuations? What about welfare or revenue in the *unknown* valuation setting, for which we obtain no nontrivial guarantees?

More generally, we leave open structural questions in constrained signalling problems more generally. We showed that there always exists a deterministic welfare-maximizing signaling scheme. What about the relative power of deterministic and randomized signaling schemes for *revenue maximization*? It was shown in [EFG⁺12] that in unconstrained settings, there always exists a revenue-maximizing scheme with at least half the optimal welfare. Does such a tradeoff hold in constrained signaling problems?

Finally, there are many other natural signaling problems one may consider. For example, what if products are given as points in a high dimensional hypercube or a high dimensional euclidean space, and an auctioneer must signal a subset of the coordinates? Problems of this form appear related to deep questions in learning theory, such as learning k -juntas and others.

References

- [ABK11] F. Armisen, C. Brownstein, and J. Krisel. Farm. *Portlandia*, 2011.
- [Ach01] D. Achlioptas. Database-friendly random projections. In *Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, page 281. ACM, 2001.
- [AHK12] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing*, 8:121–164, 2012.
- [Ake70] G.A. Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84:488–500, 1970.

- [CW09] K.L. Clarkson and D.P. Woodruff. Numerical linear algebra in the streaming model. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 205–214. ACM, 2009.
- [DNS05] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*, pages 610–618, 2005.
- [DS06] Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the 17th ACM Symposium on Discrete Algorithms (SODA)*, pages 1064–1073, 2006.
- [EFG⁺12] Y. Emek, M. Feldman, I. Gamzu, R. Paes-Leme, and M. Tennenholtz. Signaling schemes for revenue maximization. In *ACM Conference on Electronic Commerce (EC)*, 2012.
- [Fei98] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- [Fei06] Uriel Feige. On maximizing welfare where the utility functions are subadditive. In *Proceedings of the 37th ACM Symposium on Theory of Computing (STOC)*, pages 122–142, 2006.
- [GNS07] A. Ghosh, H. Nazerzadeh, and M. Sundararajan. Computing optimal bundles for sponsored search. In *Workshop on Internet and Network Economics (WINE)*, 2007.
- [GRU12] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. In Ronald Cramer, editor, *TCC*, volume 7194 of *Lecture Notes in Computer Science*, pages 339–356. Springer, 2012.
- [HR10] Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *FOCS*, pages 61–70. IEEE Computer Society, 2010.
- [KN10] D.M. Kane and J. Nelson. A derandomized sparse johnson-lindenstrauss transform. *Arxiv preprint arXiv:1006.3585*, 2010.
- [LM10] J. Levin and P. Milgrom. Online advertising: Heterogeneity and conflation in market design. *American Economic Review*, 100:603–607, 2010.
- [Mil10] P. Milgrom. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70:62–70, 2010.
- [MS12] P.B. Miltersen and O. Sheffet. Send mixed signals – earn more, work less. In *Workshop on Internet and Network Economics (WINE)*, 2012.
- [MW82] P. Milgrom and R.J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.
- [Oxl92] J. G. Oxley. *Matroid Theory*. Oxford University Press, 1992.
- [RR10] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In Leonard J. Schulman, editor, *STOC*, pages 765–774. ACM, 2010.

- [Von08] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 39th ACM Symposium on Theory of Computing (STOC)*, pages 67–74, 2008.